

UNITED STATES OF AMERICA

TITLE: **METHOD AND APPARATUS FOR PACKET ORDERING IN A
DATA PROCESSING SYSTEM**

5

FIELD OF THE INVENTION

This invention relates to data processing systems, and more particularly to a method and system for ordering packets in a multi-processor system.

10 **BACKGROUND OF THE INVENTION**

Network processing at multi-gigabit data rates, for example at 100-192 or higher data rates, requires multiple multi-threaded processors. The number of processors in a multi-processor system is limited by current integrated circuit technology. Network processing at multi-gigabit data rates requires packet
15 buffering to be done internal to the network processor. The amount of embedded memory is also limited by current integrated circuit technology. In order to properly process multiple packets in a multi-processor system, strict packet ordering between the incoming and outgoing packet path must be maintained. The problem is to maximize the number of processors and
20 minimize the number of packet buffers required while ensuring strict packet order.

A number of approaches to this problem have been attempted in the art. One approach involves removing packets from the processors in the order of
25 completion. The packets are buffered until processing of the earlier packets is completed. This approach suffers from a number of drawbacks, which include increased internal memory requirements, increased routing resource requirements, and additional operations to move data.

30 A second approach known in the art involves allowing packets to remain in processor memory until processing of the earlier packets is completed. This approach also suffers from a number of drawbacks which include increased internal memory requirements, increased packet routing resource requirements, and the problem of processor stalling and/or thread stalling
35 while waiting for the earlier packets to be processed.

Accordingly, there remains a need for a solution, which addresses the shortcomings and improves on the known approaches.

SUMMARY OF THE INVENTION

5 The present invention provides a method and system for packet ordering in a multi-processor data processing system.

10 According to one aspect of the invention, an ordering buffer is provided to maintain strict packet order in an environment where packets are not necessarily processed in order, and the buffering of packets occurs in on-chip processor memory. The ordering buffer contains a pointer and completion flag for each packet being processed or already processed but not released for output. The ordering buffer allows packet data to be read from the processor memory regardless of the completion order of processing the packet. A
15 packet is released for output in order when the processing of earlier packets has been completed.

20 Advantageously, processing of subsequent packets continues even if the processing of an earlier packet has not completed. The number of packets that can be processed ahead of an earlier packet is only limited by the number of entries in the ordering buffer.

25 The present invention provides an approach, which does not require additional memory to buffer completed packets while waiting for an earlier packet to complete.

30 In a first aspect, the present invention provides a system for processing multiple incoming data packets and outgoing data packets in a multi-processor data processing system, the system comprises: (a) means for inputting each of the incoming data packets in a specific order and means for
35 assigning an ordering pointer to each of the packets of data, the ordering pointers being stored in an ordering buffer; (b) means for processing the incoming data packets; (c) means for setting a completion flag upon completion of processing of the associated incoming packet, and said completion flag being stored in said ordering buffer with the ordering pointer associated with said incoming data packet; (d) means for outputting the

processed data packets after the associated completion flags have been set, the means for outputting being responsive to the ordering pointers associated with the incoming data packets so that the specific order of the incoming packets is maintained.

5

In another aspect, the present invention provides a method for processing multiple incoming data packets and outgoing packets in a multi-processor data processing system, the method comprises the steps of: (a) inputting each of the incoming data packets in a specific order and assigning
10 an ordering pointer; (b) processing each of the incoming data packets; (c) setting a completion flag for each of the incoming data packets upon completion of processing of the associated incoming packet; (d) outputting the processed incoming data packets after the associated completion flags have been set, the processed incoming packets being outputted based on the
15 ordering pointers associated with the incoming packets so that the specific order is maintained.

In a further aspect, the present invention provides a network processor for processing multiple incoming data packets and outgoing packets in a data
20 processing system, the system comprises: (a) an input component for inputting each of the incoming data packets in a specific order and a component for assigning an ordering pointer to each of the incoming data packets, the ordering pointers being stored in an ordering buffer; (b) one or more processor components for processing the incoming data packets; (c) a
25 component for setting a completion flag upon completion of processing of the associated incoming packet, and the completion flag being stored in the ordering buffer with the ordering pointer associated with the incoming data packet; (d) an output component for outputting the processed incoming
30 packets after the associated completion flags have been set, the output component being responsive to the ordering pointers associated with the incoming packets so that the specific order of the incoming packets is maintained for the output.

Other aspects and features of the present invention will become
35 apparent to those ordinarily skilled in the art upon review of the following

description of specific embodiments of the invention in conjunction with the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

5 Reference will now be made to the accompanying drawings which show, by way of example, a preferred embodiment of the present invention, and in which:

10 Fig. 1 shows in block diagram form a multi-processor network processor according to the present invention;

 Fig. 2 shows in diagrammatic form operation of a distributor control module in a multi-processor environment according to the present invention;

15 Fig. 3 shows in diagrammatic form operation of a collector control module in a multi-processor environment according to the present invention;

 Fig. 4 shows in diagrammatic form operation of an ordering module in a multi-processor environment according to the present invention.

20

DESCRIPTION OF THE PREFERRED EMBODIMENT

 As shown in the accompanying Figs. 1 to 4, the system according to the present invention is directed to a multi-processor system.

25 Reference is first made to Fig. 1, which shows in diagrammatic form a multi-processor or pipeline network processor according to one aspect of the present invention. The network processor is indicated generally by reference 10. As shown in Fig. 1, the network processor 10 receives incoming packets from an input device 2. The network processor 10 processes the incoming
30 packets (i.e. data) and outputs outgoing packets, which are transmitted to an output device denoted generally by reference 4. The network processor 10 finds widespread application as will be apparent to those skilled in the art. For example, the incoming device 2 may comprise a POS-PHY physical device or HDIC controller and the outgoing device 4 may comprise a router switch
35 fabric. In another application for the processor 10, the incoming device 2

comprises a router switch fabric and the outgoing device 4 comprises a POS-PHY physical layer device.

As shown in Fig. 1, the network processor 10 according to the invention
5 comprises an input control and input queue module 12, a packet memory module 14, a distributor module 16, a series of packet processors (1 to N) denoted generally by reference 18, a collector module 20, an ordering module 22, and an output control and output queue module 24. The ordering module 22 includes an ordering buffer 26 (Fig.2) according to the invention. The
10 ordering module 22 controls the operation of the ordering buffer 26 as described in more detail below.

The ordering buffer 26 as shown in Fig. 2 comprises a contiguous number of memory locations or registers 28, shown individually as 28a, 28b to
15 28m. Each of the registers 28 stores a pointer which references the location of the data packet in the packet memory 14. Each of the memory locations or registers 28 also includes a register 30 for storing a complete flag which is associated with the data packet referenced by the pointer.

20 The pointers are written into the ordering buffer 26 and the complete flag is set in the order the processing of data packets is completed by individual packet processors 18. The location 28 of the pointer in the ordering buffer 26 is based on a sequence number. The distributor module 16 assigns a sequence number to the data packet when the packet is de-queued from
25 the incoming queue or buffer 12 (i.e. by the distributor module 16). The pointers stored in the ordering buffer 26 are then en-queued onto the outgoing queue or buffer 24 in sequential order after the complete flag is set for the associated data packet. It will be understood that each sequence number corresponds to a single entry in the ordering buffer 26, and that a sequence
30 number can only be used by one data packet at a time. When the incoming data packet is de-queued from the input queue 12, the packet is assigned a sequence number. When the processed data packet is en-queued onto the output queue or buffer 24, the sequence number is released into a pool of unassigned sequence numbers.

The operation of the network processor 10 with the ordering module 24 and the ordering buffer 26 is now described with reference to Figs. 2 to 4. In Figs. 2 to 4, the network processor 10 is depicted, and the operation, described in terms of a queuing layer interface 100, a dispatch layer 200, a processing layer interface 300, and a memory interface 400.

The queuing layer interface 100 comprises an input data packet pointer queue 102 and output packet pointer queue 104.

10 The dispatch layer 200 in the network processor 10 comprises the ordering buffer 26 (as described above), a packet parsing module 202, a distributed load balancing module 204, a collection load balancing module 206, and a sequence number module 208.

15 The processing layer interface 300 comprises the packet processors 18a to 18N. As shown each of the packet processors 18 comprises a processor packet memory 302, a scheduled pointer queue 304, a free pointer queue 306, and a completed pointer queue 308.

20 The memory interface 400 comprises the packet memory 14. The interface 400 also includes a packet structure memory 402.

Reference is made to Fig. 2, which depicts the operation of the distributor module 16 (Fig. 1), and packet data write and packet statistics write operations. The distributor module 16 de-queues the packet pointer for the data packet from the input queue or buffer 12. The distributor module 16 then assigns the data packet a sequence number and copies the data packet into the processor memory 14 for processing by the packet processors 18.

30 Referring to Fig. 2, the distributor control module 16 de-queues the packet memory pointer from the input packet pointer queue 102 as indicated by path 201. The distributor module 16 selects one of the packet processors 18 based on load balancing as determined from the load balancing module 204 as indicated by path 203. Next, the distributor module 16 de-queues a pointer for the processor packet memory 302 as indicated by path 205. The pointer is de-queued from the free pointer queue 306 in the selected packet

processor 18a. The distributor module 16 then assigns the data packet a sequence number which is obtained from the sequence number module 208 as indicated by path 207. The sequence number module 208 provides a pool of sequence numbers.

5

Next, the distributor module 16 performs a packet data write operation. The packet data write operation involves reading the data packet from the packet memory 14, as indicated by path 209, and writing the data packet into the processor packet memory 302 as indicated by path 211 in Fig. 2.

10

The distributor module 16 next performs a packet statistics write operation. The packet statistics write operation involves writing a packet memory pointer and the sequence number into one of the registers 28, 30 in the ordering buffer 26 as also indicated by path 211. The packet statistics are also written into the processor packet memory 302 as indicated by paths 211 and 213. Next, the distributor module 16 en-queues the processor packet memory pointer into the scheduled pointer queue 304, as indicated by path 215 in Fig. 2.

20

Reference is made to Fig. 3, which depicts the operation of the collector module 20 (Fig. 1), and the packet statistics read, packet data read, and packet completion indication operations. In general terms, the collector module 20 copies the processed packet data out of the processor packet memory 302 and into the packet memory 14. The collector module 20 writes the packet pointer into the ordering buffer 26 at the address 28 set by the sequence number and the complete flag is also set in the register 30.

25

Referring to Fig. 3, the collector module 20 determines the packet processor 18 by reading the collection load balancing module 206, as indicated by path 221. The collector module 20 then de-queues a pointer for the processor packet memory 302 of the selected packet processor 18N, as indicated by path 223. The pointer is de-queued from the free pointer queue 306 in the selected packet processor 18N.

35

For the packet statistics read operation, the collector module 20 reads the packet memory pointer, the sequence number, and a DMA (Direct

Memory Access) command from the processor packet memory 302 as indicated by path 225 in Fig. 3.

For the packet data read operation, the collector module 20 transfers
5 the packet data from the processor packet memory 302 in the packet processor 18N to the packet memory 14, as indicated by paths 225 and 227 in Fig. 3.

For the packet completion indication operation, the collector module 20
10 first writes the packet memory pointer for the data packet into Register 28 in the ordering buffer-26 which is indexed by the sequence number as indicated by path 229. The sequence number was assigned to the data packet (as described above for Fig. 2). Next, the collector module 20 en-queues the freed pointer for the processor packet memory 302 on the free pointer queue
15 306, as indicated by path 231.

Reference is next made to Fig. 4, which depicts the operation of the ordering module 22 (Fig. 1), and the packet egress ordering operation. As will be described in more detail, the ordering module 22 walks the ordering buffer
20 26 in sequence. The ordering module 22 en-queues a packet pointer onto the output queue 104 only if the complete flag has been set. Once a packet pointer is en-queued, the ordering module 22 clears the complete flag, releases the sequence number for use by another incoming packet and the complete flag for the next entry is tested. When the ordering module 22
25 completes the last entry in the ordering buffer 26, the ordering module 22 moves back to the first entry in the ordering buffer 26 and the process is repeated.

Referring to Fig. 4, the ordering module 22 first increments an internal
30 addressing counter and waits for the complete flag to be set in the register 30 in the ordering buffer 26, as indicated by path 241. Next, the ordering module 22 en-queues the pointer (i.e. the packet pointer) for the data packet on the output packet pointer queue 104, as indicated by path 243. The ordering module 22 also clear the complete flag in the register 30 of the ordering buffer
35 26. The ordering module 22 then returns the sequence number used for this

data packet to the sequence number module 208, as indicated by path 245 in Fig. 4.

The present invention may be embodied in other specific forms without departing from the spirit or essential characteristics thereof. Certain adaptations and modifications of the invention will be obvious to those skilled in the art. Therefore, the presently discussed embodiments are considered to be illustrative and not restrictive, the scope of the invention being indicated by the appended claims rather than the foregoing description, and all changes which come within the meaning and range of equivalency of the claims are therefore intended to be embraced therein.

[illegible]